



The Impact of IA-64 on the SAS System for Decision Support

Information for Software Developers and IT Managers

Version 1.0 10/25/99

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

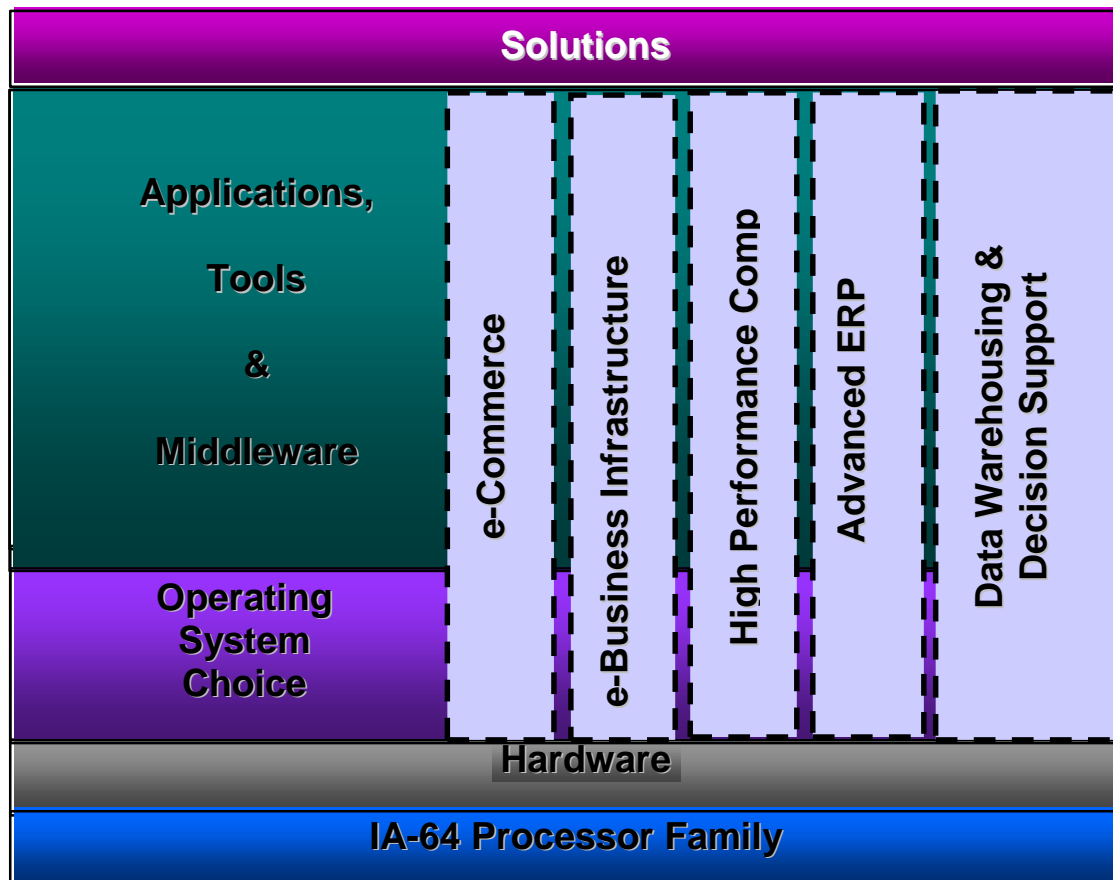
The Intel Itanium™ processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

*Copyright © Intel Corporation 1999 *Third party brands and names are the property of their respective owners.*



October 1999

With the depth of experience gained in the 16-and 32-bit worlds, Intel is now targeting its advanced IA-64 architecture at business-critical enterprise computing solutions (see figure). The first Intel 64-bit product, the Intel® Itanium™ processor, will appear in the second half of 2000 with capabilities aimed at enterprise computing. Intel, with the IA-64 architecture and the Itanium™ processor, is delivering a hardware platform designed to offer significant improvements in scalability, availability and overall application performance. The end result is not simply a faster processor.



Intel is working in close concert with SAS Institute, the world's leader in integrated data warehousing and decision support solutions for the enterprise. SAS Institute customers, who have come to depend on SAS® software for turning massive amounts of data into business intelligence, will do so much more rapidly and effectively with systems powered by the IA-64 architecture. Thus the ultimate payback to customers will be the competitive advantage gained from having the fastest-possible access to this information and knowledge, and then using that knowledge to reach critical business decisions.

Commitment to Optimization

SAS Institute has a long-standing development relationship with Intel, having been an early tester and adopter of a number of Intel's 32-bit processors, including the Pentium® and Pentium Xeon™ processors. Specifically, the latest releases of the SAS data warehousing and decision support software have been optimized for the latest generation of 32-bit Intel processors, delivering robust, high-performance solutions for today's demanding enterprise computing requirements. SAS Institute's track record of exploiting the versatility and power of Intel's 32-bit servers (<http://www.sas.com/partners/intel/backgr.html>) indicates how thoroughly SAS software will be optimized for the IA-64 architecture.

October 1999

SAS software and IA-64

SAS Institute's work is now proceeding on the optimization of SAS software for Intel's IA-64 family of products, including the Itanium™ processor. This effort will provide new levels of performance and scalability for data warehousing and decision support implementations in networked and distributed computing environments. In addition, SAS Institute customers who have benefited from Intel-based 32-bit servers will be assured that those platforms and the new IA-64 platforms will coexist.

"Working closely with Intel to understand how to prepare for and take advantage of the new architecture was a critical first step in porting our existing 32-bit applications to IA-64," said Gary Mehler, SAS Institute's manager of PC research and development. "With Intel's continued training and support, we were able to demonstrate SAS software running on a simulated IA-64 server more than a year ago. Our efforts since then have been to continue development of our application code as well as to assist Intel in improving the simulator environment for other software vendors."

Will the features planned for the IA-64 architecture score big with users? They will if they meet the needs of today's complex computing environments. Consider the US Census Bureau in Washington, D.C. The Bureau is actually an organization with more than 30 separate divisions, each of which can and usually does specify its own platforms. It's about as hybrid a computing environment as you'll find anywhere.

The Bureau handles a lot more than the census. It is the place where numbers are crunched and statistics are revealed on diverse data like housing starts, crime and balance of trade. The Bureau has used SAS software extensively since the late 1980s. Robert Bateman, chief of the SAS branch at the Bureau, says that the ability of SAS software to work with data from multiple platforms is one big reason why SAS software use is so widespread there. Bateman also applauds the "legendary analytical capabilities" of the SAS solutions.

The Importance of Balanced Systems

Looking at computing platforms, Bateman maintains that pure horsepower is secondary to what he termed "true balanced systems" – those which offer hearty processor speed along with strong I/O throughput. "What good is speed if your I/O bandwidth is tied up?" Bateman asks. "You have to get lots of data off secondary storage and into processor memory very quickly."

Bateman notes further that for his and other hybrid computing environments, a key processor attribute will be the ability to host multiple operating systems. Finally, Bateman contends that a processor architecture is only as good as the system utilities provided by the operating system and third-party vendors that support it. Intel, Bateman observes, "has always enjoyed a great deal of third party support, so I'd say the software will be abundant." And of course with Bateman's work being for the Census Bureau, scalability is of obvious importance.

A look at some of the features of the IA-64 architecture show that there is plenty of reason for the buzz surrounding this widely anticipated, next generation enterprise hardware platform. These features also address the wishes and concerns of users like Bateman who know how critical a scalable, high-performance, balanced system is to overall application performance and throughput, and therefore real-world results.

IA-64 Benefits to SAS Solutions

Large memory addressability. For the enterprise customer, having large physical memory simply translates into the ability to access more data more quickly from system memory, since much more data can be held near in the CPU for faster calculations and data analysis. This will be key to improving both operating system performance and application performance, notes Leigh Ihnen of SAS Institute's Applications Division. Ihnen says that large memory addressability will allow larger file system caches for read ahead and write behind I/O operations and also allow SAS applications to

October 1999

retain large amounts of data in memory instead of repeatedly reading the data from disk. Ultimately, this feature will help deliver decision support data extracted from warehouses much more quickly. In SAS data mining applications, for example, analyses on very large volumes of data can be performed. In the past, this analysis might have been impractical from a time perspective.

Multiple OS support. The IA-64 architecture will support all major operating systems, a key requirement in the enterprise computing environment today due to the hybrid nature of most corporate computing systems. This is also key for SAS Institute customers, who have long appreciated the ability of SAS software to run in a variety of OS environments and to access and analyze data in almost any format. As one large enterprise user said of SAS software, “It is our universal decoder ring.” And within the context of the IA-64 architecture, SAS software will run that much more efficiently in all major OS environments, including Windows NT, Monterey (the joint effort by IBM, SCO and Sequent), Solaris, and HP-UX.

Parallelism support on multiple levels--Instruction level parallelism (ILP) is the ability to execute multiple bundles (three instructions in a bundle) at the same time. The IA-64 architecture can deliver faster performance by executing multiple bundles per clock cycle. The IA-64 architecture, both at the instruction level and at the SMP system level permits more efficient use of virtually all system resources to allow and enable improved scalability. Scalability is critical to data warehouse architects, who have found through experience that many data warehouses may start out modest in size and usage. However, once the decision support benefits of analyzing clean, warehoused data become known to strategists and planners, the demands on data warehouses can skyrocket. Thus, warehouses should be planned from the ground-up with platforms that can scale to meet this demand, lest performance lag and, eventually, usage, as today’s users want their information and analysis in Internet time. As this need increases, IA-64 based systems will be able to handle higher processing loads due to their inherent support for parallelism.

The combination of IA-64 instruction set features such as **Speculation and Predication, along with a rotating register model and a large number of registers** enable compilers to vastly improve the optimization for performance of the complex analytical sub-routines found through out decision support and data mining applications.

Speculation and predication reduce the memory latency during data access and improve the efficiency of branch handling. These features are particularly vital in calculation-intensive applications that manipulate large volumes of data, such as SAS decision support applications that model customer behavior for inclusion in Customer Relationship Management (CRM) systems. Predication removes branches through parallel execution to increase performance. Speculation reduces impact of memory latency, thus providing significant benefits to applications with many cache accesses. When used in a large-scale decision support system, these features allow the processor to achieve higher throughput since it will be able to obtain data from memory more efficiently as well as run program code in a more predictable – and therefore faster- way.

Large number of registers and rotating registers. These features of the IA-64 architecture allow more data to be handled within the registers, the fastest storage of all in a computer, making for a more flexible and faster applications environment. For a complex application like the SAS System, which processes vast amounts of data, these additional registers will translate into higher throughput rates. Meanwhile, the rotating registers further enhance loop performance with an innovation called software pipelining, which avoids code expansion and thus reduces cache misses. When software pipelining is used effectively in a large application like the SAS System, processor resources such as registers and logic units can be kept more uniformly busy, resulting in higher throughput.

When all of these factors are taken into account, the bottom line for SAS customers is that IA-64-based servers will be built to scale, store, process and access larger amounts of data to yield significant performance increases of data and compute-intensive applications. That’s just what SAS decision support applications are: sophisticated, analytical tools that harvest, crunch through, and convert volumes of data into business intelligence and knowledge. The SAS System

October 1999

provides an environment enabling data warehousing, business intelligence, and customer relationship management applications, each of which has arrived at the vanguard of contemporary enterprise computing.

About SAS Institute

For more than 20 years now, SAS Institute has developed and marketed the world's foremost decision support solutions. In fact, SAS software is at work delivering critical business information to managers and strategists in 98 of the Fortune 100 companies (and nearly 90 percent of the Fortune 500), and is used in total by 3.5 million users in 115 countries. A hallmark of SAS software historically and today is its ability to collect data from nearly any platform and data format; to cleanse and then transform that data into information by providing multidimensional analysis tools that are both client/server and Web enabled. SAS Institute continues to lead the software industry in R&D outlays, committing nearly one of every three dollars in revenue to product and solution development. For more information visit SAS Institute's Intel Alliance web site at <http://www.sas.com/partners/intel>.

About Intel

Intel, the world's largest chip maker, is also a leading manufacturer of computer, networking and communications products. Additional information about Intel is available at www.intel.com/pressroom.